Evolution and complexity

Miconi, T.

DOI:

[10.1162/artl.2008.14.3.14307](https://doi.org/10.1162/artl.2008.14.3.14307)

Document Version

Publisher's PDF, also known as Version of record

Citation for published version (Harvard):

Miconi, T 2008, 'Evolution and complexity: The double-edged sword', *Artificial Life*, vol. 14, no. 3, pp. 325-344.
<https://doi.org/10.1162/artl.2008.14.3.14307>

[Link to publication on Research at Birmingham portal](#)

Publisher Rights Statement:

Thomas Miconi, Massachusetts Institute of Technology

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Evolution and Complexity: The Double-Edged Sword

Thomas Miconi*
University of Birmingham

Keywords

Evolution, complexity, arrow of complexity, trends

Abstract We attempt to provide a comprehensive answer to the question of whether, and when, an *arrow of complexity* emerges in Darwinian evolution. We note that this expression can be interpreted in different ways, including a passive, incidental growth, or a pervasive bias towards complexification. We argue at length that an arrow of complexity does indeed occur in evolution, which can be most reasonably interpreted as the result of a passive trend rather than a driven one. What, then, is the role of evolution in the creation of this trend, and under which conditions will it emerge? In the later sections of this article we point out that when certain proper conditions (which we attempt to formulate in a concise form) are met, Darwinian evolution predictably creates a sustained trend of increase in maximum complexity (that is, an arrow of complexity) that would not be possible without it; but if they are not, evolution will not only fail to produce an arrow of complexity, but may actually prevent any increase in complexity altogether. We conclude that, with regard to the growth of complexity, evolution is very much a double-edged sword.

I Introduction: The Arrow-of-Complexity Hypothesis

It appears that the maximum level of organismal complexity in the biosphere has increased, in some sense and at a certain scale, over time. Bacteria are simpler than humans, if only because humans contain millions of (interacting) cells, each much more complex than bacteria. Sponges are simpler than flatworms, which are simpler than trilobites, which are simpler than birds. This crude, large-scale growth of the maximum of complexity seems to hold for any notion of organismal complexity we know of (and, as we will see, there are quite a few of them). To our knowledge, no serious author has explicitly denied an overall increase in the maximum of complexity at a global scale, though some have urged caution [35]. At a slightly lower scale, there is some evidence that the maximum number of cell types (a common measure of organismal complexity) has increased over time within metazoans [6, 53].

Faced with such a remarkable increase in complexity, some authors have suggested a fundamental tendency of evolution to create more and more complexity. This position actually predates Darwin, being found in Lamarck. One variant of this notion is the *arrow-of-complexity* hypothesis [4]:

[The] hypothesis of the arrow of complexity asserts that the complex functional organization of the most complex products of open-ended evolutionary systems has a general tendency to increase with time.

* School of Computer Science, University of Birmingham, Edgbaston, Birmingham B152TT, UK. E-mail: t.miconi@cs.bham.ac.uk

It turns out that this apparently straightforward statement has generated significant amounts of confusion and misunderstanding. We need not go very far to find one example: In the same article, the author goes on to express that this hypothesis is not accepted by Gould [24], Maynard Smith and Szathmáry [32], and McShea [35]. Yet these authors, and those works, do *not* deny that the complexity of the most complex organism increases in evolution. At most, McShea professes “emphatic agnosticism.” Gould, however, writes (p. 169): “I do not challenge the statement that the most complex creature has tended to increase in elaboration through time”—an almost verbatim statement of the arrow-of-complexity hypothesis.¹ As for Maynard Smith and Szathmáry, their book is a description of several successive “major transitions” in the history of evolution, many of which can reasonably be interpreted as increases in complexity (more precisely, hierarchical complexity [35]).

How can such confusion arise? The crux of the matter, of course, is that while nobody seems to utterly deny the *fact* that maximum complexity has increased over the entire course of evolution, there is much debate regarding the *interpretation* of this fact: As we will see, what these authors explicitly question is not that increasingly complex individuals are found in nature, but rather the notion of a built-in evolutionary *drive*, or *bias*, towards increased complexity in evolution.

It is one of the achievements of recent evolutionary theory to have shown that this controversy results in no small part from a conjunction of misunderstandings. The crux of the matter is that the question of whether “there is a tendency for complexity to increase in evolution” is a highly ambiguous question. This is because almost every word in it can be understood in many different ways, leading to much unnecessary confusion. When fuller attention is given to details (and to the implicit assumptions that may lurk behind apparently innocuous terms), it is realized that the sustained growth of maximum or average complexity does not necessarily imply a pervasive bias towards complexity in evolution.

2 What is Complexity?

2.1 Descriptive and Functional Complexity

If we are to study the relationship between evolution and complexity, we need to explain what we call “complexity” in the first place, if only to avoid misunderstandings. It turns out that definitions and measures of complexity abound throughout the literature (for a broad panorama see [15, 35, 1]). We will concentrate specifically on two notions of complexity, which we believe are of particular interest when discussing the effects of Darwinian evolution: descriptive (or morphological) complexity, and functional complexity.

Descriptive or morphological complexity is easily defined: It is the amount of information that is necessary to describe an object. In the words of Dawkins [11]: “A crab is morphologically more complex than a millipede because, if you wrote a pair of books describing each animal down to the same level of detail, the crab book would have a higher word-count than the millipede book. The millipede book would describe a typical segment then simply add that, with listed exceptions, the other segments are the same. The crab book would require a separate chapter for each segment and would therefore have a higher information content.” However, very much in the same way as with Shannon entropy (another measure of information content), it turns out to be maximal for completely disordered objects. A pile of rubble consists of many more parts, each much more different from any other, than a standing wall of bricks, so its descriptive complexity would be higher.

What about functional complexity? We choose to simply define functional complexity as the descriptive complexity of a *functional* system. A “system” is any set of interacting parts. We call it “functional” if the interaction of the parts produces a certain external result, which would not be

¹ In a different, related piece [23], Gould writes: “No one can doubt that more complex creatures arose sequentially after this prokaryotic beginning—first eukaryotic cells, perhaps about two billion years ago, then multicellular animals about 600 million years ago, with a relay of highest complexity among animals passing from invertebrates, to marine vertebrates and, finally (if we wish, albeit parochially, to honor neural architecture as a primary criterion), to reptiles, mammals and humans.”

obtained if some (or any) of the parts failed to behave just as it does. Note that the result must come from the actual interaction of active parts, not just from their mere existence. Functional complexity is simply descriptive complexity applied to any such functional system, that is, a growing function of the number of parts and of their mutual differences.

While we admit that this is a rather crude definition, we believe that it does capture much of the intuitive notion of functional complexity within a system. In particular, it does retain the idea of difficulty (in a design sense), or improbability, which we associate with highly complex biological features: Clearly a certain system that depends on the precise behavior of many parts is somehow more improbable or surprising than one that only depends on the collaboration of a few. Similarly, it stands to reason that systems relying on the interactions of many similar parts are somehow easier to arrive at than systems relying on the precise features of many differentiated parts.

This notion of functional complexity may be contrasted with that of Heylighen [25], which emphasizes the joint occurrence of separation and interconnection; or with that of Dawkins [10], which stresses the improbability of the result (rather than the “emerging improbability” resulting from the highly constrained interactions of many parts). For example, in our view, Rube Goldberg contraptions (phantasmagoric constructions in which highly intricate machinery is used to obtain trivial results) would be regarded as functionally complex, whereas Dawkins’ view would regard them as comparatively noncomplex, due to their underwhelming output. Functional complexity, in our sense, is more about the functioning than the function being performed.

2.2 Implications of the Definitions

With this notion of functional complexity in mind, we can already make a few inferences. The most important one is that adaptive functional complex systems are *rare* in the space of all possible systems. Any given set of parts is unlikely to be actually performing any meaningful function, let alone an adaptive one. This unlikelihood increases with the descriptive complexity of the system: The more parts there are, the more improbable it is that they will harmoniously cooperate (a requirement for functional complexity). This is important, because it implies that generating random systems would be a very inefficient way to come up with functionally complex systems, while other search methods (such as Darwinian evolution) might, under the right conditions, prove more successful.

Another observation is that although adaptive functional complexity is rare in the space of all systems, it comes in *loose clusters*: coming up with a new, more complex, yet still successful design is typically much easier if we use an already complex, successful design as a starting point. More precisely, slightly altering a successful complex system has a higher chance of yielding another system of equivalent or higher complexity than random search. If we have access to the mechanism of a working clock, random alterations to this existing mechanism are much more likely to result in other working clocks than random assemblages of gears. Importantly, this is not a logical necessity; rather, it is a reasonable inference, which we may assume holds for many system spaces.

3 What is a Tendency?

While having a working notion of complexity is important, it appears that much of the confusion actually derives from the notion of *tendency*. What do we mean when we say that the maximum or average complexity “tends” to increase in evolution?

3.1 Trends, Passive Trends, and Driven Trends

McShea [34, 35, 37] has devoted much work to the conceptual clarification of trends and mechanisms in evolution. The most widely known result of this work is probably the operative distinction between passive and driven trends. A *trend* is any sustained directional change in a summary statistic of a certain quantity (such as maximum or average size, complexity, etc.) within a particular clade. Such a trend is called *driven* if it can be shown to result from a pervasive, consistent *bias*, applying to

Figure 1. Passive and driven trends. In both of these graphs, the maximum and the average follow a trend of sustained increase through time. In the left-hand graph, these trends are *passive*: The branching process is not biased towards higher or lower values. A trend in the maximum emerges mechanically by envelope expansion. A trend in the average also emerges, because the process is constrained by a hard lower bound that cannot be crossed. Both trends, resulting from a globally unbiased process, are passive. In the right-hand graph, however, the trends of increase in the average and the maximum are *driven*: The underlying branching process is clearly biased towards higher values.

all subclades; for example, if speciation or extinction events are consistently biased towards higher values of the quantity. *Passive* trends, on the other hand, are those that occur mechanically in globally unbiased processes, possibly constrained by a local inhomogeneity—for example, a *hard* boundary that the value cannot cross. Figure 1 illustrates this distinction graphically.

It is important to note that passive trends may well be enforced by active, external forces, such as natural selection. For example, we might have a situation where no member of a certain clade is selectively viable below a certain size, or complexity level. In this case, provided there is no constraint on the maximum, the average would grow mechanically over time.

Furthermore, even in a globally passive trend, the small-scale behavior of each individual lineage may well be everywhere governed by a local, driven trend—as long as these many local trends do not exhibit any consistent global direction. In other words, *passive trends may well be everywhere driven*² [34, p. 1751]. Conversely, nonselective forces may enforce driven trends. Thus the important distinction between passive and driven trends does not lie in the source or mechanism of the bias that creates this trend, but solely on whether this bias is *local* or *pervasive*, respectively, throughout the space explored by the clade [37].

3.2 Benefits of the Terminology

This sound terminology goes a long way to identify (and prevent) possible sources of misunderstandings. Restated in McShea's terminology, the arrow-of-complexity hypothesis simply states that evolution spontaneously generates a trend of increase in the maximum of organismal complexity over time. It is silent on the question of whether this trend is driven (the result of a pervasive bias in Darwinian evolution towards higher complexity) or passive (the result of a direction-blind branching process, mechanically pushing the external envelope).

We can now see where much of the controversy arises: What some authors (especially Gould [24]) reject is not the existence of the trend, but the idea that it should necessarily be driven, rather than passive. This, in turn, seems to be misconstrued by others as denying the existence of the trend itself, or the role of evolution in its emergence: The implicit assumption is that any trend must of necessity be driven, and therefore, that denying bias is tantamount to rejecting the trend itself. This seems to explain why Gould, for example, is portrayed as a “skeptic” regarding the existence of a trend [4], or even as denying that “complexity increases in evolution” [1]. Gould, as we saw, certainly

2 This point deserves particular attention: In everyday language, the words “passive” and “driven” seem to be mutually exclusive. But in McShea's terminology, a trend is passive if it is *globally* unbiased, which does not imply that its local behavior at any point must be: Even in a globally unbiased (passive) branching process, each particular branch may well be strongly biased (driven) in its own local direction. Any overall trend will still be called passive if the many local biases are uncorrelated, resulting in a globally unbiased process.

agrees that complexity increases in evolution; What he rejected was the idea of a general, pervasive bias or drive towards complexity. This illustrates the high potential for confusion that lies behind such deceptively straightforward terms as “trend” and “tendency.”

4 Reformulating the Question

Using McShea’s terminology, we can now reframe our previously ambiguous question in a more satisfactory manner. We admit that, at a sufficiently large scale, there has been a trend of increasing complexity in evolution: There has (so far) been an arrow of complexity from exclusively simple forms to a variety of simple and complex forms. This point is uncontroversial and widely accepted, though appropriate caution is necessary [35]. We are now left with two questions:

1. Has this trend been driven (whether by selection or by some other mechanism), or passive? Is there a built-in bias in evolution to push average or maximum complexity preferentially upwards, in general?
2. What is the role of the evolutionary mechanism in the emergence of this trend? What are the necessary conditions for the arrow of complexity to arise in Darwinian evolution?

The following sections will address both questions in turn.

5 Theoretical Justifications for a Driven Complexity Growth

It is interesting that, independently of the evidence for or against a general, driven trend of increasing complexity in evolution, several theoretical justifications have been suggested for such a driven trend. We will discuss a few typical examples.

5.1 The Assumption of Complexity Advantage

An apparently common (though often implicit) viewpoint is that functional complexity should be favored by selection, in general, because living organisms can generally be made more efficient by increasing their functional complexity. A statement of this position is apparently offered by Darwin himself, in the crucial fourth chapter of the *Origins of Species* [9] (quoted with this intent by Rosslenbroich [45]):

If we take as the standard of high organisation, the amount of differentiation and specialisation of the several organs in each being when adult ... , natural selection clearly leads towards this standard: for all physiologists admit that the specialisation of organs, inasmuch as in this state they perform their functions better, is an advantage to each being; and hence the accumulation of variations tending towards specialisation is within the scope of natural selection.

However, the continuation of this passage (*not* quoted by Rosslenbroich) paints an altogether different picture:

On the other hand, we can see ... that it is quite possible for natural selection gradually to fit a being to a situation in which several organs would be superfluous or useless: in such cases there would be retrogression in the scale of organisation....

But it may be objected that if all organic beings thus tend to rise in the scale, how is it that throughout the world a multitude of the lowest forms still exist; and how is it that in each great class

some forms are far more highly developed than others? Why have not the more highly developed forms everywhere supplanted and exterminated the lower? On our theory the continued existence of lowly organisms offers no difficulty; for natural selection, or the survival of the fittest, does not necessarily include progressive development—it only takes advantage of such variations as arise and are beneficial to each creature under its complex relations of life. And it may be asked what advantage, as far as we can see, would it be to an infusorian animalcule—to an intestinal worm—or even to an earthworm, to be highly organised. If it were no advantage, these forms would be left, by natural selection, unimproved or but little improved, and might remain for indefinite ages in their present lowly condition. And geology tells us that some of the lowest forms, such as the infusoria and rhizopods, have remained for an enormous period in nearly their present state.

To the naive view that natural selection spontaneously leads to complexification, Darwin opposes a more neutral notion of evolutionary progress, according to which species simply adapt to their own particular conditions. This may, or may not, lead to increases (or decreases) in complexity over time. In other words, the flaw of the complexity-advantage assumption is that, while it rightly recognizes the possible advantages of complexity, it fails to consider the possible advantages of simplicity. The world is sufficiently complex that there are an almost infinite number of ways to “earn a living,” and on the whole, evolving lineages may just as readily gain advantage through simplicity as through complexity.

5.2 Modern Endorsements of the Complexity-Advantage Assumption

Despite Darwin’s own objections, several authors have used the complexity-advantage assumption as an implicit justification for the idea that Darwinian evolution by natural selection entails a drive toward complexity.

5.2.1 Coevolution and Arms Races

One way in which evolution is sometimes said to favor the increase of complexity is through coevolution between lineages that must adapt to each other’s evolving features. This is expected to lead to an *arms race* [12], in which coevolving lineages should come up with more and more adaptive features to defeat each other. The argument is summarized by Rosin and Belew [44]:

When the parasites evolve to overcome this failure [against the hosts], they create new challenges for the hosts; the continuation of this may lead to an evolutionary “arms race” [10]. New genotypes arise to defeat old ones. New parasite types should serve as a drive toward further innovation, creating ever-greater levels of complexity and performance by forcing hosts to respond to a wider range of more challenging parasite test cases.

It is interesting that these authors, like several others, cite Dawkins as a source. Dawkins, coauthor of the original arms-race hypothesis, has explicitly and emphatically denied that the progress brought about by arms races was intrinsically linked to complexification—as we shall see in Section 5.4.

But at any rate, the problem with this position is easily identified: It implicitly assumes that the most straightforward way to cope with an opponent’s increase in complexity is a corresponding, countering increase in complexity. In other words, complexity calls for complexity, because only complexity can cope with complexity. Unfortunately, as many practitioners of artificial coevolution (including Rosin and Belew themselves) have remarked [55, 17, 38], the opposite is often true: Simplicity is often quite efficient at dealing with complexity, with devastating effects. The most obvious example is the phenomenon of *opportunism*, or parasitism, whereby promising complex organisms are systematically defeated by simplistic opponents which successfully exploit some of their weak points (what Watson and Pollack call “focusing on the wrong thing” [55]). We may also mention *mediocre stable states* [17] (also called “collusion” [42] or “immature inductions” [3]), where

competitors remain stuck in a small, uninteresting section of the search space from which none has any incentive to budge—the coevolutionary equivalent of a local optimum. As a special case of the latter, we may mention the possibility of cycles [39], in which competitors cycle through a short, predictable sequence of simple behaviors without any long-term gain in general performance or complexity.³

In general, these problems can be summarized by this simple observation: Evolution is interested in immediate advantage, rather than in complexity or generality, and there seems to be no general relationship between the two: To reliably obtain the latter, one must forcibly connect it to the former by means of algorithmic contrivance [39, 44, 50, 48, 38].

To be sure, there are many examples of specific situations where certain coevolutionary pressures have clearly driven increases in complexity; some will be discussed in Section 6. What must be doubted is not the existence of coevolution-driven increases in complexity, but that these should be a general rule rather than interesting specific cases.

5.2.2 A Cybernetic Argument for the Complexity Advantage

Heylighen [25] uses a cybernetic argument to justify the complexity advantage assumption. In brief, natural selection should favor organisms that are able to cope with a larger variety of perturbations; but, by application of Ashby's law (Heylighen argues), this in turn requires higher countering variety in the behavior of these organisms, and thus more functional complexity:

[Ashby's law of requisite variety] states that in order to achieve control, the variety of actions a control system is able to execute must be at least as great as the variety of environmental perturbations that need to be compensated. The larger the variety of available counteractions, the larger the set of disturbances that can be corrected, and the larger the domain of potential environmental situations in which the control system can survive. All other things being equal, greater control variety implies greater fitness.... Therefore, evolution through natural selection will tend to increase control, and thus, because of Ashby's law, internal variety. This can be interpreted as a functional differentiation, i.e., the appearance of more diverse activities or functions....

All other things being equal, a system that can survive situations A, B, and C, is absolutely fitter than a system that can only survive A and B. Such an increase in absolute fitness is necessarily accompanied by an increase in functional complexity. Thus, evolution will tend to irreversibly produce increases of functional complexity.

This leads to interesting questions. For example, by any measure of organismal complexity, beetles are arguably less complex than chimpanzees. Yet there do not seem to be many "situations" that could be survived by chimpanzees, but not by beetles—as opposed to the other way round. More generally, it is not immediately clear that more complex individuals can cope with more environments than simpler individuals.

Quite possibly, the crucial difficulty in this argument is that Ashby's law applies when one needs to actively counter certain variations in order to homeostatically control certain variables; but in nature, *adapting* to a change in an external variable does not necessarily entail *countering* this change. Evolving lineages may well cut their dependency on certain variables (as in previously carnivorous lineages turning into herbivores), or simply ignore the variation in question. To use Heylighen's own example of warm-bloodedness, we note that birds and mammals are the only surviving lineages of warm-blooded animals: Insects, whose range, diversity, and total biomass vastly exceed that of either mammals or birds, have hardly suffered from their lack of internal temperature control mechanisms.

3 Interestingly, examples of such coevolutionary cycles have actually been reported in nature [49].

If nature is any guide, surviving more “situations” can be obtained by simplicity just as readily as by complexity, which contradicts the premise of the argument.⁴

To summarize, it seems that the cybernetic argument, like the arms race argument, attributes a general advantage to complexity by disregarding the potential advantages of simplicity: It essentially assumes that only complexity (in the evolving lineage) can deal with complexity (in the environment).

5.3 Cope’s Rule and the Size-Complexity Insight

Bonner [6] has proposed a way in which natural selection might indirectly drive an increase in the maximum of complexity. The basis of this argument is Cope’s rule, that is, the generalization that there is a tendency for lineages to grow in size over time. Bonner provides a justification for this rule by pointing out that “there is always room at the top”: “by becoming larger the organisms enter new size worlds where, among other things, they avoid predation and competition.” Bonner also notices that size is (very roughly) correlated with complexity: On the whole, larger organisms tend to be more complex. Bonner takes the practical approach of defining complexity as the number of cell types in an organism.

Bonner argues that the relationship between size and complexity is explained by considerations of engineering and efficiency: Creating and maintaining a larger body is likely to require higher division of labor. Thus, a selective drive for size creates an indirect selective drive towards complexity, though Bonner is very clear that this is only one component among many of selective action on complexity (p. 221). The book repeatedly points out that this selection-driven increase in size and complexity only applies at maximum values: Intermediate forms have certainly been maintained, and “every biological community on the globe today contains a complete spectrum from the smallest to the largest organism” (p. 221). In addition Bonner emphasizes that selection for increase and decrease in size and complexity have occurred concomitantly throughout evolution.

This argument is not without difficulties. First, it relies on Cope’s rule, which is itself the subject of recurrent controversy within evolutionary research [26, 2, 22]. But more generally, the argument does not seem to provide any compelling evidence to show that the observed trend is indeed driven (biased), rather than passive (diffusive). In fact, the very possibility of a passive trend is simply not considered at all. It should be pointed out that the book was written before McShea and Gould began to draw attention to the question of trend mechanisms. Because the possibility of a passive trend is not considered, the increase in maximum value is seen as requiring an explanation, a causative force, which happens to be the “room at the top” principle.⁵

In fact, with hindsight, the wealth of data and evidence provided by Bonner appears to be not only compatible with, but strikingly suggestive of, a passive trend. This applies particularly to ancestry data taken from the fossil record (such as Figures 8 and 9 in the book), which do not seem to show any pervasive upward bias, but rather an apparently undirected branching process constrained by a lower bound. Also, Gingerich’s evidence [21, 20], which Bonner discusses several times (including in the passage quoted above), is suggestive of a passive trend: When no pervasive bias exists, but rather species evolve both up and down indifferently along the scale and the sole constraint is a local boundary, we may expect that the overall rate of change will decrease as we average it over longer periods.⁶

5.4 Expert Opinion Against the Complexity-Advantage Assumption

More generally, theoretical justifications of a complexity advantage face a difficulty: The variety of environments (both biotic and abiotic) faced by evolving lineages, as well as their remarkable

4 The author extends the cybernetic argument to coevolutionary situations and arms races, which can be answered in a similar manner: Adapting to an “increase in variety” in a competing species does not entail matching this variety. Adaptation is simply not a problem of control in Ashby’s sense.

5 Incidentally, this may show that Gould’s and McShea’s argument is not quite as obvious or routine as it may have seemed to some [11].

6 This interpretation of Gingerich’s data now seems consensual: See Weiner [56], quoted in pp. 106–109 of [7]. Valentine and colleagues [53] provide refined data on the increase of cell types among metazoans, and model this increase with a passive system.

plasticity, implies that at any time there is a very wide range of ways to “earn a living” in nature. Optimizing the efficiency of a population in its chosen way of life (or, if the opportunity arises, switching trades altogether) is what Darwinian evolution does. There does not seem to be any compelling reason why the majority of natural lifestyles should benefit from continuous increases in complexity.

This point has been admirably expressed by Dawkins, in an unduly overlooked piece [11] from which we will now quote at length:

Why should any thoughtful Darwinian have expected a majority of lineages to increase in anatomical complexity? Certainly it is not clear that anybody inspired by adaptationist philosophy would.... Our human line happens to have specialised in complexity, especially of the nervous system, so it is only human that we should define progress as an increase in complexity or in braininess. Other species will see it differently, as Julian Huxley... pointed out in a piece of doggerel entitled “Progress”:—

The Crab to Cancer junior gave advice:
Know what you want, my son, and then proceed
Directly sideways. God has thus decreed
Progress is lateral; let that suffice.

Darwinian Tapeworms on the other hand
Agree that Progress is a loss of brain,
And all that makes it hard for worms to attain
The true Nirvana peptic, pure and grand.

Man too enjoys to omphaloscopize.
Himself as Navel of the Universe...

... Huxley's tapeworms, using a parasite-centred definition of progress, see the point with opposite sign [from human authors, who focus on increasing complexity]. A statistically minded swift would search in vain for evidence that a majority of evolutionary lineages show trends towards improved flying performance. Learned elephants, to borrow a pleasantry from Steven Pinker..., would ruefully fail to uphold the comforting notion that progress, defined as a driven elongation of the nose, is manifested by a statistical majority of animal lineages.

In this quotation, the opinions of three authors (all of them well-known evolutionists, all of them staunch supporters of the power of natural selection) are cited. It is interesting that this passage is taken from a review of Gould's own book on the subject of progress [24], which on the question of complexity adopts pretty much the same position. In the beginning of this section we showed that Darwin had already expressed very similar feelings. Let us also cite Maynard Smith and Szathmáry, who unambiguously reject the notion that evolution should be biased towards complexity in the first chapters of their book [32] (this chapter carries the rather explicit title “The fallacy of progress”). Similarly we may quote Mayr [33]:

Among all the forms of life, neither the prokaryotes nor protists, fungi or plants has evolved smartness, as it should have if it were “better.” In the 28 plus phyla of animals, intelligence evolved in only one (chordates) and doubtfully also in the cephalopods. And in the thousands of subdivisions of the chordates, high intelligence developed in only one, the primates, and even there only in one small subdivision. So much for the putative inevitability of the development of high intelligence because “it is better to be smart.”

All in all, we note that there is much expert opinion against the idea that natural selection should drive a pervasive trend towards increasing functional or behavioral complexity in evolution. Rather, the consensus seems to point towards Darwin's original conception of evolution by natural selection: a recklessly opportunistic process, which will constantly exploit any possible way to make a living, independently of whether it leads to increases or decreases in complexity.

6 The Evidence: Complexity in Evolution

6.1 The Enduring Dominance of Simple Life Forms

A significant argument against the idea of an active, pervasive, selection-driven bias towards complexity growth over the course of evolution is that simpler life forms still dominate the planet. In a typically engaging section, Gould [24] famously argued that “on any possible, reasonable or fair criterion, bacteria are—and always have been—the dominant forms of life on Earth.” Prokaryotes (Bacteria and Archaea) are more ancient, more numerous, and more diverse, have conquered more environments, and apparently compose vastly more of the cellular biomass than all other living entities combined (a position only reinforced by later estimates [57]). Simpler prokaryotes have clearly not been displaced by their more complex eukaryotic cousins.

We note that it is possible to extend this pattern to other taxonomic levels. The most functionally complex animals, besides humans, are arguably great apes—of which, it turns out, many happen to be on the verge of extinction. We may compare this with the success (whether in numbers, range, or conservation prospects) of “lesser”⁷ simians, especially the genus of macaques, the most widespread of all primate genera. Nonhuman primates in general are significantly *less* successful, by any measure of success, than the placid artiodactyls (ruminants, pigs, etc.). Chordates (the phylum that includes vertebrates) are vastly exceeded in number, diversity, and total biomass by arthropods, the truly dominant form of metazoan life [5]. Land animals in general are simply overwhelmed by land plants in terms of total biomass. The dominance of prokaryotes, Gould's “modal bacter,” merely crowns what appears to be a general pattern, or rather the absence of a pattern: At any given taxonomic level, more functionally complex entities are simply not more successful than their simpler counterparts. In fact, a lot could be said for the opposite position, though of course counterexamples abound.

The important fact here is not the enduring existence of some primitive life forms, *per se*. What is significant is that at almost any given scale, and for any subtree of the Tree of Life, we observe that simpler organisms not only endure, but often seem to dominate the more complex organisms. This leads to the following question: If evolution is indeed generally biased towards complexity, then why do only a minority of life forms become more complex? Why did all those simple forms fail to feel the supposed pull of the complexity drive? This pattern would suggest that those taxa that did reach high complexity did so because of certain particular, specific selective pressures, to which they somehow adapted by increasing their complexity.

6.2 Examples of Particular Selection-Driven Trends Towards Complexity

Considering the amount of confusion that has been pointed out by McShea, Gould, Dawkins, and others, it may not be entirely futile to stress the obvious: Nobody denies that in many cases, active, selection-driven increases in complexity do in fact occur. Complex adaptations often do confer a definite selective advantage over relatives that do not possess them. The result is that they will mechanically propagate throughout the population. Such trends have happened at all levels, ranging from individual subspecies to the entire biosphere, both within and across lineages. We will only give a few examples, which we believe are of particular interest.

⁷ Quantitative measures of behavioral complexity and neocortex development among primates are given by Dunbar [14].

For a specific example, we need look no further than the evolution of our own species, *Homo sapiens* [7, 29]. Modern humans are the latest stage in a sequence of various groups, starting with the early australopithecines (essentially bipedal apes), going through the early *Homo* branch (the habilis group: *H. rudolfensis*, *H. habilis*, etc.), the erectus group (*H. ergaster*, *H. erectus*, etc.), the archaic *Homo sapiens* (*H. Heidelbergensis*, *H. Neanderthalensis*), and eventually modern humans, presumably descended from the same branch as *H. Heidelbergensis* (in parallel with the Neanderthal derivation). Now this sequence shows two notable features. First, each new group showed demonstrably more complex behavior than its predecessors, as can be deduced from artifacts. A trend of increase in brain size is also apparent (especially in certain stages, such as the erectus and sapiens groups [29]). Second, at a sufficiently large scale, these groups tended to thrive while their predecessors were displaced. It is therefore reasonable to argue that this sequence constitutes a trend, driven by natural selection, whereby increasingly complex groups successively displace their predecessors. Is it possible that the specific case of human evolution, in which natural selection did drive a sequential increase in complexity, influenced our view of evolution in general?

However, even in this most indisputable of trends, it is important to exert caution. First of all, the appearance of early *Homo* certainly did not displace the australopithecines. On the contrary, australopithecines continued to expand and radiate, with a trend towards robust forms (*A. boisei*, *A. robustus*), for at least another million years [19]. The divergence somewhere along the line of the early *Homo* branch (the habilis group) was apparently a “much more isolated and insignificant trend” [29, p. 328], which locally reversed the general tendency. Only with the erectus group did *Homo* finally achieve prominence. Furthermore, it is not absolutely clear that the displacement of each group by its successor was invariably caused by the latter’s higher complexity: Climate changes have clearly had an influence on the evolution [28] (and apparently the extinction [18]) of various human lineages. Thus, even in the clearly driven trend of human complexification, the pluralism of evolution shows its signs.

More generally, the evolution of complex organs constitutes a remarkable, ubiquitous example of selection-driven increase in complexity. Historically, the paradigm of complex organ evolution is the eye, a high feat of engineering that has been discovered many times in the course of evolution. Dawkins [10] has famously described the highly complex sonar of certain bats, using it as a replacement for the often-used vertebrate eye as an example of complex organ evolution. The evolution of such complex organs, each composed of several highly heterogeneous, interdependent parts delicately assembled to obtain a certain result (and thus functionally complex according to our definition), has necessarily involved a long sequence of gradual increases in complexity, each of them preserved by selection. In the case of the eye, many of these intermediate stages have actually been found in living species [43, Figure 13.2]. Of course, we should keep in mind that many creatures (perhaps most impressively ants) do very well with low-fidelity eyes.

The massively enlarged primate brain must of necessity be the result of a selection-driven trend; this is because primate brains are costly structures, especially in terms of energy. As Dunbar [14] explains, “because the cost of maintaining a large brain is so great, it is intrinsically unlikely that large brains will evolve just because they can.” Dunbar provides a wealth of evidence that indicates that the relative size of the neocortex in primates (including humans) correlates remarkably with measures of social behavior (such as the size of the social group), but not with ecological aspects such as diet, foraging behavior, or range size. This would suggest that social pressures, rather than ecological factors, were the proximal drive for larger neocortices in primates—a clear example of a coevolution-driven complexity increase.

The evolution of complex chemical cycles within cells may be seen as a counterpart (and predecessor) to the evolution of complex organs in multicellular organisms. Perhaps the most remarkable case is aerobic respiration, an intricate oxygen-based cycle that extracts energy from complex organic molecules. This highly efficient process, yielding up to 36 ATP molecules for each glucose molecule (against two for anaerobic fermentation), is now extremely widespread among living creatures: Exclusively anaerobic creatures, once dominant, have been relegated to a secondary role (though by no means eliminated altogether) by creatures endowed with the more complex, but more

powerful aerobic respiration apparatus.⁸ Of course, anaerobic fermentation is still present as a component of the aerobic respiratory cycle, to the effect that aerobic cells may revert to anaerobic fermentation in the absence of oxygen (much to the satisfaction of certain higher mammals not entirely averse to the effects of fermented beverages).

More generally, it is widely acknowledged that the general architecture of modern living systems (based on nucleic acids and proteins) is too complex to have arisen spontaneously, in its present form, from abiotic reactions. Therefore, intermediate forms (now extinct) must have preceded the appearance of modern life. The utter elimination of these early systems, which we may surmise to have been quite varied, can hardly be explained by pure contingency. In fact, it is reasonable to posit that these ur-organisms were actually displaced by modern life, which simply proved more efficient in assimilating precious carbon atoms and free energy. From this viewpoint, the total domination of modern life can be seen as a selection-driven trend, through which a sequence of increasingly complex systems successively displaced their simpler, less efficient ancestors.⁹ This process, however, has not endured into the evolution of modern life, as can be shown by the massive dominance of prokaryotic forms in the present biosphere.

6.3 Examples of Particular Selection-Driven Decreases in Complexity

Having shown particular examples of driven trends towards complexity, we feel obliged to emphasize that trends often run in the other direction as well.

One obvious form of selection towards lower complexity is the emergence of parasitism in previously free-living species: An organism that was originally able to survive on its own finds itself dependent on its host. This often leads to a reduction in complexity, as Gould [24, p. 220] predictably observes.¹⁰ This phenomenon is well known and does not require further emphasis.

Another common source of selectively driven decrease in complexity is miniaturization. Bonner [6], pursuing his argument that size and complexity often go hand in hand, mentions the case of dwarf rotifers, which, because of new biomechanical constraints imposed by their nanism, replace their gut system with a continuous cytoplasm nested with small food-carrying vacuoles. The result is a decrease both in structure differentiation and in number of cell types. A well-known example of miniaturization (and associated loss of complexity) is that of Acari, particularly among the most minute mites. For example, dust mites have shed off much of their digestive structure (including their stomach), and as a result must ingest the same piece of food several times and use external digestion.

More generally, in parallel with the constant elaboration of complex organs, an equally constant process of complexity reduction occurs in nature. We need only cite a few well-known examples. The most obvious telltales of evolutionary simplification are vestigial structures. The vermiform appendix in humans is the remnant of a structure that allows its bearer to digest cellulose in plants, something that humans cannot do: This is clearly a diminution in functional complexity, by any measure. The small pink pouch in the inner corner of our eyes is the nonfunctional remnant of the nictitating membrane (the “third eyelid”) found in many other animals, including cats and dogs. Modern whales retain small, undifferentiated bone structures as a memory of what used to be a perfectly functional pelvis and hind legs [51]. This also affects the paradigmatic eye: In many creatures that have reverted to dark environments, the functional complexity of their visual systems has decayed into nonfunctionality, sometimes culminating in the complete elimination of the eye itself (the cavefish *Astyanax* is interesting in that various stages of eye loss can be observed in distinct modern

8 Even then, things are not quite clear-cut: The fermenters' higher rate of consumption may actually allow them to overcome respirators in certain situation of direct topical competition [41].

9 In fact, Fernando and Rowe [16] suggest that the very origin of life itself is the result of such a selection-driven trend towards increasing complexity, initiated by natural selection among prebiotic autocatalytic systems.

10 However, Gould also reminds us that in many case, parasites may also exhibit high complexity at some stages of their life: *Saculina* is mentioned as an example both of extreme simplicity (in its almost amorphous host-bound form) and complexity (in its remarkably intricate development cycle).

populations [27]). Even the pinnacle of animal complexity, the vertebrate brain, has undergone significant morphological and functional simplification in certain lineages: Let us mention salamanders [46] and certain bats [47].¹¹

Finally, we will mention a rather spectacular form of successful complexity decrease. This example, tentatively suggested for this purpose by McShea [36] (and recently confirmed by further studies [13]), concerns the appearance of infectious tumor cells. In at least two separate examples (one affecting dogs in a rather benign way, the other affecting Tasmanian devils with a usually fatal outcome), it has been found that a certain lineage of tumor cells could be transmitted from host to host through contact, with the result that the tumors effectively behaved as populations of single-cell parasites of the host species. In both cases, the tumor cells were shown to share a common origin (being genetically similar to each other across hosts, and different from their hosts), and thus can be considered a distinct species. For our purpose, the key point is that these single-cell species originated from mutated cells within multicellular, “higher” organisms. In other words, in at least two separate occurrences, a *major transition* of evolution [32] (namely, the invention of multicellularity) was successfully reversed. This dramatic drop in complexity from highly complex mammals to single-cell organisms created new, simpler, yet successful species, especially in the case of the dog tumor lineage, which is thought to be at least several centuries old.

7 Conclusion: The Passive Trend of Complexity Growth in Evolution

7.1 Passive Growth of Complexity

In summary, both theory and evidence suggest the same interpretation: There has been a trend of increase in the maximum of complexity in evolution, and this trend seems to have been resolutely passive—that is, globally unbiased. At all scales, individual lineages will be driven to whatever form of adaptation evolution may encounter, which may or may not lead to increases, or decreases, in complexity. There does not seem to be any overarching bias, either in evidence or in theory, that would drive a majority of lineages towards higher complexity. The Tree of Life does not “seek” complexity as actual plants seek sunlight.

Again, it is important to note that this does *not* contradict the hypothesis of an arrow of complexity; it does, however, contradict certain interpretations, namely those that surmise a built-in tendency for Darwinian evolution to drive lineages towards increasing complexity. This does not imply a rejection of the existence of a trend: A passive trend is still a trend.

7.2 Nontriviality of Passive Trends

In addition, we must stress that a passive trend does not in any way imply that Darwinian evolution is neutral or irrelevant to complexity growth. There is a tendency to see passive trends as somehow trivial, or not requiring any particular explanation. That this is not the case is of fundamental importance. Complexity has indeed risen to extraordinary levels, and the ease with which nature comes up with complex forms demands explanation.

The difficulty of obtaining a sustained arrow of complexity is amply demonstrated by the field of artificial evolution, in which thorough use of Darwinian evolution has indeed created appreciable levels of complexity [48, 30], but nothing that could even be compared to the enormous growth of complexity in nature. In fact, the difficulty of initiating open-ended evolution in artificial experiments is widely noted [31]. Consequently, the very fact that evolution seems unbiased towards higher or lower complexity is in itself worthy of note, because in artificial settings the process seems very much biased *against* continuing complexity growth. The apparent absence of a bias for or against complexity in nature only seems trivial if we forget how strong the bias against complexity seems to be in computers.

¹¹ Brain evolution in vertebrates is a complicated business, despite a highly conserved general structure. Many fishes have similar brain size to mammals and birds, relative to their body size. Many independent increases in brain size have been identified, but stability and decreases have also occurred [40].

7.3 New Questions

This casts a new light on our problem. If a passive trend has occurred in the evolution of complexity, the question then becomes: What makes such a trend *possible*? What conditions are necessary for the appearance and sustained increase of complexity (that is, for the emergence of an arrow of complexity) in a system? What is the role of the Darwinian mechanism in this emergence? This is the second part of our reformulated question (as put forth in Section 4), with which the remaining sections of this article will be concerned.

8 The Role of Darwinian Evolution in Complexity Growth

8.1 Systems, Designs, and Complexity

In this section we want to evaluate the effects of Darwinian evolution on complexity, in the most abstract, general sense. Our perspective is unabashedly conceptual: We want to gain insight into the relationship between evolution and complexity, rather than devise technical methods to be used in actual experiments. To do this, we ask which constraints on complexity exist in abstract nonevolutionary generative systems; then we ask how the introduction of an evolutionary mechanism alters these constraints.

Let us consider an abstract generative system, of which we only know two things:

- A certain *design generator* is being used to generate new designs; the particulars of the design generator define a certain space of possible designs, which can be explored.
- New designs are being generated according to a certain heuristic; in other words, the space of available designs is being sampled in a certain sequence, as specified by the heuristic.

For now, we assume that the heuristic used to generate new designs is unknown: It could be random, evolutionary, hill-climbing, or anything else. We now ask the question: What is necessary for such a system to produce unbounded growth in functional complexity?

A moment's thought suggests that in principle, at least two (rather obvious) conditions are absolutely necessary. First, the system generator must in principle be able to generate arbitrarily complex designs. Second, the search heuristic must not impose a hard upper bound on adaptivity or functional complexity of the designs being sampled. It need not favor complexity—it may even discriminate against it, but not to the point of forbidding it entirely above a certain level. Provided these two conditions are met, over time the system should mechanically produce arbitrarily complex designs, regardless of which heuristic is being used.

8.2 Limitations of Nonevolutionary Complexity Growth

8.2.1 Inefficiency

However, another moment's thought will make it clear that, without any further device, such a system might well take impossibly long times to come up with any such complexity. This is because, as we have seen in Section 2.2, functionally complex designs are rare in the space of all possible designs. A random design generator may come up with arbitrarily complex designs, given enough time, but what good would that be if “enough time” exceeds the thermodynamic limit of universal heat death?

8.2.2 Finite Populations and Finite Lifetimes

Another difficulty occurs when designs are eventually removed from the system. For example, designs might have a specified maximum lifetime (which would mean that all designs are eventually eliminated), or the population could have a maximum size (which implies that some designs must be eliminated).

In this case, while the system will still come up with designs of unbounded functional complexity over time, we have no assurance that the maximum of complexity at any time will *sustainedly* increase. This is because, without any further assumption, we have no way to know that the rate at which

complex designs are being generated balances the rate at which they will be removed: In fact, the opposite is much more likely. While arbitrarily complex designs will eventually occur, these designs may very well be mere blips in an otherwise consistently low background: If the current maximum of complexity is removed, we have no assurance that it will be replaced by designs of similar complexity. Again, this is because functionally complex designs are rare within the space of all possible designs. The consequence is that, barring further assumptions, the population may well spend most of its time at a very low level of functional complexity.

8.2.3 Consequence: No Arrow of Complexity

We must therefore conclude that, while the two obvious requirements expressed above are sufficient to generate arbitrary complexity over (much) time, they are not sufficient to create what we would call an arrow of complexity. This would imply a *sustained, noticeable* increase in the maximum complexity, rather than the highly occasional appearance of fleeting, isolated complexity spikes.

8.3 Effect of Evolution on Complexity

8.3.1 The Darwinian Heuristic

Now let us introduce a new requirement: We posit that our system is evolutionary, that is, it follows the *Darwinian heuristic* as elegantly formulated by Wagner and Altenberg [54]: “the Darwinian heuristic is to choose sample points by perturbing the more fit ones among those sampled thus far.” In other words, new designs are not generated arbitrarily: Rather, existing designs are evaluated after a certain metric (their *fitness*), and those designs that exhibit a relatively high fitness are selected as starting points for small jumps towards newer designs. Our question is: How does this new specification affect the emergence of functional complexity?

8.3.2 The Darwinian Mechanism Restricts Exploration

The first thing to note is that the Darwinian heuristic actually imposes a *restriction* on the search process: The search process will now follow certain directions at the expense of others. This restriction constrains both the scope and the direction of the search process. The restriction in scope comes from the fact that, because evolution proceeds by small jumps from existing designs, newer designs can only be chosen in the neighborhood of the current population. The restriction in direction comes from the fact that only those designs that exhibit a high fitness will be used as starting points for further exploration.

In unfavorable conditions, this restriction may severely limit the exploration of design space. For example, if the population happens to be in a certain domain of the search space, which is entirely surrounded by relatively unsuccessful designs, then the process may well end up being trapped in that domain. Conversely, potentially highly fit designs may never be reached if they are surrounded by a region of unfit designs. This suggests at least two conditions for complexity to appear in evolutionary systems, namely, that it should be *successful* and *reachable*.¹²

8.3.3 Benefits of the Restriction

On the other hand, this restriction has important potential benefits. First of all, when successful designs are found, they will tend to be preserved, even beyond the lifetime of individuals—or, at the very least, they will stand an immensely higher chance of being preserved than under random search, as long as they remain successful within the context of the current population. This permits the persistence of functional complexity, despite its rarity in design space, even in the face of finite lifetimes: Complexity is no longer condemned to occur only in short blips over an otherwise simple background.

¹² In the context of this discussion, a design is *successful* if its fitness is high enough, relative to the current population, to be used as a starting point for further samples. A point in design space is *reachable* if certain successful designs within the current population might be altered into the corresponding design.

Furthermore, as we have already noted in Section 2.2, adaptive functional complexity tends to come in loose clusters in design space: If we want to obtain a successful complex design, we stand a much higher chance to obtain it by slightly altering another successful design of lower complexity than by randomly poking at the entire design space. The consequence is that, by preserving existing (successful) complexity, the Darwinian heuristic also facilitates further increases in complexity, provided such increases are possible within the neighborhood of the population: A successful complex design may act as a *stepping stone* towards even higher complexity. This is, in essence, the basis of the *blind watchmaker* mechanism, so vividly expounded by Dawkins [10]: By slowly accumulating small beneficial improvements, each working as a stepping stone towards the next, the Darwinian heuristic may eventually produce enormous amounts of adaptive functional complexity, with incomparably higher probability than under random search.

Thus we see that certain requirements must hold if evolution is to generate designs of unbounded complexity: It is necessary that, at all times, successful designs of higher complexity exist within the current neighborhood of the population. At first sight, these conditions might seem rather straightforward. However, are they really sufficient?

8.3.4 Enter the Red Queen: Variable Fitness and Complexity

So far we have not considered a certain aspect of fitness, namely the fact that it is not necessarily stable over time. A design that is successful at a certain time may later prove unsuccessful if conditions have changed. This is particularly the case in situations of coevolution, that is, when the fitness of a given design depends at least in part on the features of other existing designs within the population. In this case, species evolve constantly in order to adapt to each other's evolution—and those that fail to do so disappear. This is the well-known Red Queen effect [52, 8].

For our purposes, the basic consequence is that even if all previously mentioned conditions are met, an evolutionary process may still fail to produce a sustained increase in complexity. Even if, at any time, there exists potentially successful complex designs in the reachable neighborhood of the current population, complexity will not necessarily grow over time. This may seem counter-intuitive: At first sight, if at any time successful complex designs exist within the reachable neighborhood, then we would expect that eventually evolution would find them, and thus complexity would grow.

The reason why fitness instability may prevent such continuous growth is that it may disrupt the stepping-stones effect on which this growth relies. If fitness is variable, then designs of intermediate complexity may become unsuccessful before they have a chance to be altered into more complex designs, putting the latter beyond reach of the process. Thus, despite the constant presence of reachable, successful designs of higher complexity, no continuous process of complexification may occur. If moderately complex designs are not sufficiently robust to serve as stepping stones towards higher complexity, the blind watchmaker fumbles, and the arrow of complexity is broken. The effects of this phenomenon in artificial coevolution have been mentioned in Section 5.2.1.

Thus we see that for complexity to appear in an evolutionary system, it not only needs to be successful and reachable: It also needs to be *robust*, even in the face of possible instabilities (including new competitors, parasites, etc.)—at least sufficiently so to provide a stepping stone for further complexity growth. This is an additional requirement that, if not observed, may actually prevent the rise of complexity in evolutionary systems.

8.3.5 Putting It All Together

How can we express these requirements as concisely as possible? We suggest the following formulation:

- At any time, there are more possible jumps towards successful designs of higher and equal complexity than towards designs that would make the starting points of these jumps unsuccessful.

Here “higher and equal complexity” is to be understood in comparison with the current maximum of the population. A jump is the creation of a new design, using an already existing design (the *starting point* of the jump) as a parent. Clearly, by definition, the starting point of a jump must be successful.

This condition is clearly statistical: Even if such jumps exist and constitute the majority of possible jumps, we have no guarantee that they will be taken in the next step, so the corresponding complexity increase may not occur immediately. Conversely, even if jumps towards higher complexity are rare, we may still occasionally encounter one of them. What this condition determines, however, is the presence or absence of a *trend*—that is, a long-term tendency. If this condition does not hold over time, then statistically we do not expect any long-term trend towards higher complexity to emerge; similarly, if it does hold consistently, then a long-term trend towards higher complexity is statistically expected.

More importantly, this condition is very much local (both in time and in design space). If the condition holds at a certain time, that does not mean in any way that it will keep doing so in the future: Any new jump runs the risk of breaking the condition. For example, one particularly unlucky jump might discover a niche of devastatingly efficient simple individuals (in the extreme, we might conjure images of “gray goo” bacteria). Should this happen, then the probability of jumps might become highly biased towards these simple individuals, and thus the condition would not hold any more. Conversely, when the condition does not hold, one particularly lucky jump may stumble upon a new area of robust complex adaptations, making the condition valid. Therefore, maximum complexity will tend to increase as long as the condition is consistently met—but only that long. Complexity will not tend to increase (or may even tend to decrease) if the condition consistently fails to hold.

8.4 Evolution and Complexity: A Double-Edged Sword

In summary, we suggest that the conditions under which Darwinian evolution will generate a sustained trend of increase in the maximum of complexity (that is, an arrow of complexity) can be summarized as concisely as possible in the following manner:

1. Arbitrarily complex adaptive designs must exist in the design space explorable by the design generator.
2. The search heuristic must not impose a hard upper bound on the functional complexity of designs.
3. At any time, there must be more possible jumps towards successful designs of higher and equal complexity, than towards designs that would make the starting points of these jumps unsuccessful.

Thus we see that there is a deep ambiguity in the relationship between evolution and complexity. This ambiguity is captured in the third requirement above, which (as we saw) is very much specific to Darwinian evolution. Random search, for example, may create arbitrary amounts of complexity over (much) time, even if this requirement does not hold, while Darwinian evolution cannot. But if (and as long as) this requirement holds, then Darwinian evolution can actually create an arrow of complexity (a sustained, noticeable trend of increase in maximum complexity), which is utterly beyond the reach of random search. In this sense, with regard to complexity, Darwinian evolution is very much a double-edged sword.

9 Conclusion

Clearly, at this point, we have gone a long way to clarify the statement that open-ended evolutionary processes have an inherent, lawlike tendency to create creatures with increasingly complicated

functional organization. We have seen that the hypothesis of an arrow of complexity can be interpreted in several ways, with different support for each interpretation. It was shown that the undeniable increase of maximum complexity in evolution is best explained as a passive (unbiased) trend, that is, a side consequence of an otherwise complexity-neutral process of reckless local adaptation.

We have also seen that Darwinian evolution is capable of creating an arrow of complexity (that is, a sustained, noticeable trend of increase in maximum complexity) that could not exist without it, provided certain requirements hold. These requirements are that at any time, higher complexity should be reachable, successful, and reasonably robust, in the sense made explicit by our formulation. They result from the restrictions imposed by the Darwinian heuristic on the exploration of search space. As long as these evolution-specific requirements are consistently met, then evolution will keep producing designs of increasing adaptive complexity in a sustained manner. But if they cease to be, then evolution will not only fail to produce an arrow of complexity—it may actually prevent complexity increases altogether. In other words, with regard to complexity growth, evolution is a double-edged sword, which may slash both ways.

Thus the appearance and growth of complexity in evolution results neither from a built-in, relentless march to “progress,” nor from the miraculous conjunction of a series of vastly improbable coincidences. Rather, it is simply ascribable to the presence of certain objective conditions, applying to the capacities of the design generator and the structure of design space. No grand scheme of cosmological destiny, no poignant tale of triumph against all odds; just the pedestrian, plainly logical, yet amazingly powerful consequence of the properties of modern life in the face of the challenges and opportunities posed by the laws of physics, the geological environment—and itself.

References

1. Adami, C. (2002). What is complexity? *BioEssays*, 24, 1085–1094.
2. Alroy, J. (1998). Cope’s rule and the dynamics of body mass evolution in North American fossil mammals. *Science*, 280, 731–734.
3. Angeline, P. J., & Pollack, J. B. (1993). Competitive environments evolve better solutions for complex tasks. In S. Forrest (Ed.), *Proceedings of the Fifth International Conference on Genetic Algorithms* (pp. 264–270). San Mateo, CA: Morgan Kaufmann.
4. Bedau, M. (1998). Four puzzles about life. *Artificial Life*, 4, 125–140.
5. Black, S., Shepard, M., & Allen, M. (2001). Endangered invertebrates: The case for greater attention to invertebrate conservation. *Endangered Species UPDATE*, 18, 42–50.
6. Bonner, J. (1988). *The evolution of complexity by means of natural selection*. Princeton, NJ: Princeton University Press.
7. Boyd, R., & Silk, J. (2000). *How humans evolved*. New York: Norton.
8. Cliff, D., & Miller, G. F. (1995). Tracking the Red Queen: Measurements of adaptive progress in co-evolutionary simulations. In *Proceedings of the Third European Conference on Artificial Life (ECAL-95)* (pp. 200–218). Berlin: Springer.
9. Darwin, C. (1859). *On the origin of species* (1st ed.). London: John Murray.
10. Dawkins, R. (1986). *The blind watchmaker*. New York: Penguin.
11. Dawkins, R. (1997). Human chauvinism. *Evolution*, 51(3), 1015–1020.
12. Dawkins, R., & Krebs, J. R. (1979). Arms races between and within species. *Proceedings of the Royal Society of London, Series B*, 205, 489–511.
13. Dingli, D., & Nowak, M. A. (2006). Cancer biology: Infectious tumour cells. *Nature*, 443, 35–36.
14. Dunbar, R. (1998). The social brain hypothesis. *Evolutionary Anthropology*, 6, 178–190.
15. Feldman, D., & Crutchfield, J. (1998). Measures of statistical complexity: Why? *Physics Letters A*, 238, 244–252.
16. Fernando, C., & Rowe, J. (2007). Natural selection in chemical evolution. *Journal of Theoretical Biology*, 247(1), 152–167.

17. Ficici, S. G., & Pollack, J. B. (1998). Challenges in coevolutionary learning: arms-race dynamics, open-endedness, and medicocre stable states. In *Artificial Life VI: Proceedings of the Sixth International Workshop on the Synthesis and Simulation of Living Systems* (pp. 238–247). Cambridge, MA: MIT Press.
18. Foley, R. (1994). Speciation, extinction, and climatic change in hominid evolution. *Journal of Human Evolution*, 26, 277–289.
19. Foley, R. (1998). Pattern and process in hominid evolution. In J. Bintliff (Ed.), *Structure and contingency: Evolutionary processes in life and human society*. Leicester University Press, Leicester, UK.
20. Gingerich, P. (1974). Stratigraphic record of Early Eocene Hyopsodus and the geometry of mammalian phylogeny. *Nature*, 248, 107–109.
21. Gingerich, P. (1983). Rates of evolution: Effects of time and temporal scaling. *Science*, 222, 159–161.
22. Gould, G., & MacFadden, B. (2004). Chapter 17: Gigantism, dwarfism, and Cope's rule: Nothing in evolution makes sense without a phylogeny. *Bulletin of the American Museum of Natural History*, 285, 219–237.
23. Gould, S. (1994). The evolution of life on earth. *Scientific American*, 271, 84–91.
24. Gould, S. J. (1996). *Full house/Life's grandeur*. New York: Jonathan Cape.
25. Heylighen, F. (1999). The growth of structural and functional complexity during evolution. In F. Heylighen & D. Aerts (Eds.), *The evolution of complexity* (pp. 17–44). Dordrecht, The Netherlands: Kluwer.
26. Jablonski, D. (1997). Body-size evolution in Cretaceous molluscs and the status of Cope's rule. *Nature*, 385, 250–252.
27. Jeffery, W., & Martasian, D. (1998). Evolution of eye regression in the cavefish *Astyanax*: Apoptosis and the Pax-6 gene. *Integrative and Comparative Biology*, 38, 685.
28. Lahr, M., & Foley, R. (1998). Towards a theory of modern human origins: Geography, demography, and diversity in recent human evolution. *Yearbook of Physical Anthropology*, 41, 137–176.
29. Lewin, R., & Foley, R. (2004). *Principles of human evolution*. New York: Blackwell.
30. Lipson, H., & Pollack, J. (2000). Automatic design and manufacture of artificial lifeforms. *Nature*, 406, 974–978.
31. Maley, C. C. (1999). Four steps toward open-ended evolution. In W. Banzhaf, J. Daida, A. E. Eiben, M. H. Garzon, V. Honavar, M. Jakicla, & R. E. Smith (Eds.), *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 99)* (pp. 1336–1343). San Mateo, CA: Morgan Kaufmann.
32. Maynard Smith, J., & Szathmary, E. (1995). *The major transitions in evolution*. Oxford, UK: Oxford University Press.
33. Mayr, E. (1995). Answer to “The abundance of life-bearing planets.” *Bioastronomy News*, 7(4).
34. McShea, D. (1994). Mechanisms of large-scale evolutionary trends. *Evolution*, 48, 1747–1763.
35. McShea, D. (1996). Perspective: Metazoan complexity and evolution: Is there a trend? *Evolution*, 50, 477–492.
36. McShea, D. (2001). The minor transitions in hierarchical evolution and the question of a directional bias. *Journal of Evolutionary Biology*, 14, 502–518.
37. McShea, D. (2005). The evolution of complexity without natural selection, a possible large-scale trend of the fourth kind. *Paleobiology*, 31, 146–156.
38. Miconi, T., & Channon, A. (2006). The N-strikes-out algorithm: A steady-state algorithm for coevolution. In G. Yen (Ed.), *Proceedings of the IEEE Congress on Evolutionary Computation (CEC 2006)*. Piscataway, NJ: IEEE Press.
39. Nolfi, S., & Floreano, D. (1998). Coevolving predator and prey robots: Do “arms races” arise in artificial evolution? *Artificial Life*, 4, 311–335.
40. Northcutt, R. (2002). Understanding vertebrate brain evolution. *Integrative and Comparative Biology*, 42, 743.
41. Pfeiffer, T., Schuster, S., & Bonhoeffer, S. (2001). Cooperation and competition in the evolution of ATP-producing pathways. *Science*, 292, 504.
42. Pollack, J. B., Blair, A. D., & Land, M. (1997). Coevolution of a backgammon player. In C. G. Langton & K. Shimohara (Eds.), *Artificial Life V: Proceedings of the Fifth International Workshop on the Synthesis and Simulation of Living Systems*. Cambridge, MA: MIT Press.

43. Ridley, M. (1993). *Evolution*. New York: Blackwell.
44. Rosin, C. D., & Belew, R. K. (1995). Methods for competitive co-evolution: Finding opponents worth beating. In L. Eshelman (Ed.), *Proceedings of the Sixth International Conference on Genetic Algorithms* (pp. 373–380). San Mateo, CA: Morgan Kaufmann.
45. Rosslenbroich, B. (2006). The notion of progress in evolutionary biology—The unresolved problem and an empirical suggestion. *Biology and Philosophy*, 21, 41–70.
46. Roth, G., Nishikawa, K. C., & Wake, D. B. (1997). Genome size, secondary simplification, and the evolution of the brain in salamander. *Brain Behavior and Evolution*, 50, 50–59.
47. Safi, K., Seid, M., & Dechmann, D. (2005). Bigger is not always better: When brains get smaller. *Biology Letters*, 1, 283–286.
48. Sims, K. (1994). Evolving 3D morphology and behavior by competition. In R. Brooks & P. Maes (Eds.), *ALIFE IV: Proceedings of the Fourth International Workshop on Synthesis and Simulation of Living Systems* (pp. 28–39). Cambridge, MA: MIT Press.
49. Sinervo, B., & Lively, C. (1996). The rock-paper-scissors game and the evolution of alternative male strategies. *Nature*, 380, 240–243.
50. Stanley, K. O., & Miikkulainen, R. (2002). The dominance tournament method of monitoring progress in coevolution. In *GECCO 2002 Workshop Proceedings*. San Mateo, CA: Morgan Kaufman.
51. Thewissen, J. G., Cohn, M., Stevens, L., Bajpai, S., Heyning, J., & Horton, W., Jr. (2006). Developmental basis for hind-limb loss in dolphins and origin of the cetacean bodyplan. *Proceedings of the National Academy of Science of the USA*, 103(22), 8414–8418.
52. Valen, L. V. (1973). A new evolutionary law. *Evolutionary Theory*, 1, 1–30.
53. Valentine, J., Collins, A., & Meyer, C. (1994). Morphological complexity increase in metazoans. *Paleobiology*, 20, 131–142.
54. Wagner, G., & Altenberg, L. (1996). Complex adaptations and the evolution of evolvability. *Evolution*, 50, 967–976.
55. Watson, R. A., & Pollack, J. B. (2001). Coevolutionary dynamics in a minimal substrate. In L. Spector, E. D. Goodman, A. Wu, & W. B. Langdon (Eds.), *Proceedings of the GECCO 2001 Conference*. San Mateo, CA: Morgan Kaufmann.
56. Weiner, J. (1994). *The beak of the finch: A story of evolution in our time*. New York: Random House.
57. Whitman, W., Coleman, D., & Wiebe, W. (1998). Prokaryotes: The unseen majority. *Proceedings of the National Academy of Science of the USA*, 95, 6578–6583.